



Key Points:

- We assembled a labeled data set of 5 Hz Global Navigation Satellite System (GNSS) velocity time series from 77 earthquakes over nearly 20 years
- We trained a supervised random forest classifier for detecting seismic motion that outperforms existing detection methods
- Improved detection enables lightweight, high rate GNSS velocity processing to be included in operational ground motion observations

Correspondence to:

T. Dittmann,
stdi2687@colorado.edu

Citation:

Dittmann, T., Liu, Y., Morton, Y., & Mencin, D. (2022). Supervised machine learning of high rate GNSS velocities for Earthquake strong motion signals. *Journal of Geophysical Research: Solid Earth*, 127, e2022JB024854. <https://doi.org/10.1029/2022JB024854>

Received 31 MAY 2022

Accepted 14 OCT 2022

Supervised Machine Learning of High Rate GNSS Velocities for Earthquake Strong Motion Signals

T. Dittmann^{1,2} , Y. Liu¹, Y. Morton¹ , and D. Mencin² 

¹University of Colorado, Boulder, CO, USA, ²UNAVCO Inc, Boulder, CO, USA

Abstract High rate Global Navigation Satellite System (GNSS) processed time series capture a broad spectrum of earthquake strong motion signals, but experience regular sporadic noise that can be difficult to distinguish from true seismic signals. The range of possible seismic signal frequencies amidst a high, location-varying noise floor makes filtering difficult to generalize. Existing methods for automatic detection rely on external inputs to mitigate false alerts, which limit their usefulness. For these reasons, geodetic seismic signal detection makes for a compelling candidate for data-driven machine learning classification. In this study we generated high rate GNSS time differenced carrier phase (TDCP) velocity time series concurrent in space and time with expected signals from 77 earthquakes occurring over nearly 20 years. TDCP velocity processing has increased sensitivity relative to traditional geodetic displacement processing without requiring sophisticated corrections. We trained, validated and tested a random forest classifier to differentiate seismic events from noise. We find our supervised random forest classifier outperforms the existing detection methods in stand-alone mode by combining frequency and time domain features into decision criteria. The classifier achieves a 90% true positive rate of seismic event detection within the data set of events ranging from M_w 4.8–8.2, with typical detection latencies seconds behind S-wave arrivals. We conclude the performance of this model provides sufficient confidence to enable these valuable ground motion measurements to run in stand-alone mode for development of edge processing, geodetic infrastructure monitoring and inclusion in operational ground motion observations and models.

Plain Language Summary Continuously operating, high sample rate Global Navigation Satellite System (GNSS) sensors that experience ground shaking from an earthquake can provide valuable data regarding the nature of the ground motion. If this data is streamed in real-time, these observations can complement existing traditional seismic infrastructure measurements that are used for earthquake early warning or rapid ground motion assessments. However, the data from these sensors can be noisy and have non-earthquake artifacts that are difficult to tell apart from true seismic signals. In this work we used a nearly 20-years archive of high sample rate GNSS velocities occurring during known seismic events to train, validate and test a machine learning model for earthquake detection. This machine learning approach is taken from existing algorithms used for a wide variety of challenging classification problems where a label can be applied to a sample. We demonstrate that this data-driven method, without any external information, is more likely to detect these signals with less false alarms when compared to existing methods. The added confidence this algorithm provides will allow these valuable measurements to be included in operational seismic assessment and warning decision criteria.

1. Introduction

Real-time, low-latency observations of medium to great earthquake ground motions are vital to rapid hazard assessment and earthquake early warning (EEW) systems. These measurements have historically been recorded by inertial seismometers. Higher rate (≥ 1 Hz) continuous Global Navigation Satellite System (GNSS) measurements capture stronger dynamic motions and permanent displacements of propagating seismic waveforms (Larson et al., 2003; Nikolaidis et al., 2001). These geodetic strong motion measurements (Larson, 2009) will rarely clip nor require double integration that leads to magnitude saturation in the near-field of larger, destructive earthquakes common to inertial velocity sensors (Bock et al., 2004; Colombelli et al., 2013; Crowell et al., 2013). Furthermore, additional material low-latency observations densify existing seismic ground motion measurements. These observations are particularly valuable when damaging seismic events occur in sparsely instrumented regions (Grapenthin et al., 2017) or when networks or infrastructure fails.

© 2022, The Authors.

This is an open access article under the terms of the [Creative Commons Attribution License](#), which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

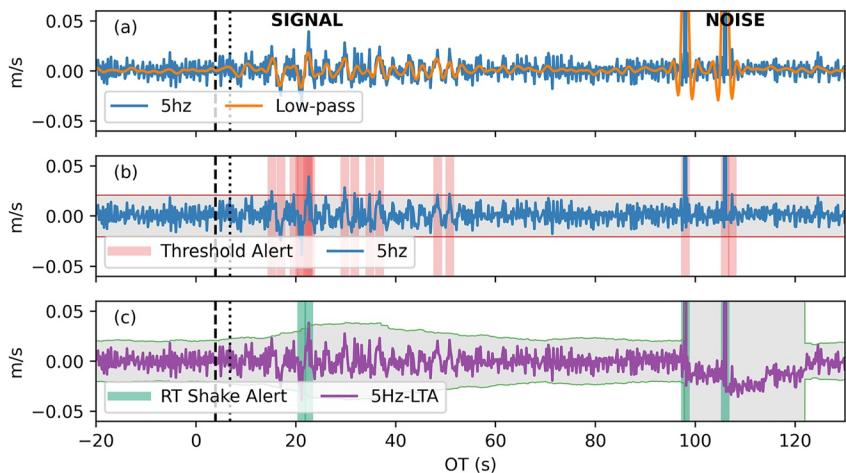


Figure 1. An example of the difficulties of differentiating a relatively weak seismic global navigation satellite system (GNSS) signal event and a GNSS noise disturbance using existing detection methods. The signal depicted is the east component of station P507 observing a M_w 5.41 (USGS event ID: ci15200401) at approx. 23 km; velocities are presented in seconds relative to the event origin time (OT). The proximal noise disturbance depicted is a non-geophysical processing artifact or signal propagation effect that might result from sources such as cycle slips, ephemeris, multipath, or other signal path effects. Panel (a) is the 5 Hz timeseries, in addition to a low-pass filtered (corner frequency of 0.5 Hz) timeseries to emphasize the signal and noise for the reader. Gray shading represent areas within the noise estimate for each respective method. Vertical dashed lines are estimated (iasp91 model) P - and S -wave arrival times. Panel (b) illustrates a static threshold taken from (Hodgkinson et al., 2020). This approach is sensitive to the weak signal, but equally sensitive to noise. The threshold has limited memory and rapidly alerts to the onset of the noise disturbance, and also issues several additional false alerts around 105 s OT. Panel (c) is a variation on an STA/LTA approach implemented from (Psimoulis et al., 2018) called RT-Shake with a moving threshold of three times the moving standard deviation. This approach detects the signal event later in the waveform with little information regarding the event duration. The noise disturbance adds an initial false alert, after which the noise region expands to minimize additional false alerts. However this memory would result in missed detection should such a noise disturbance occur immediately prior or during a seismic event.

However, geodetic deformation timeseries are noisier than traditional inertial sensors (Melgar et al., 2020). GNSS seismic waveforms are observed from spaceborne radio signals traveling over 20,000 km through often convoluted atmospheric signal paths. The radio signals are tracked by a variety of receivers using antenna situated in a range of radio frequency environments. As a result, these measurements have relatively high and complex noise signatures, making separating signal from noise challenging. Signal amplitudes from the largest, most costly events can be difficult to distinguish from non geophysical events, such as filter reconvergence or signal loss of lock (Figure 1a). Medium magnitude events, often difficult to detect above the geodetic noise floor, can be destructive or tsunamigenic. The ability to make accurate, low-latency distinction between true signals and noise in stand-alone mode, without external sensors or information, minimizes points of failure and decision latency and maximizes the value of these integral network decision inputs and potential edge processing capabilities.

Current approaches to detect motion use variations of time domain thresholds to flatten the decision to a function of signal amplitude. Several existing approaches make use of low-pass filters similar to traditional STA/LTA seismological phase picking (i.e., Allen & Ziv, 2011; Goldberg & Bock, 2017; Kawamoto et al., 2016; Ohta et al., 2012; Minson et al., 2014) that extract static offsets for finite fault inversion but filter valuable dynamics information. Recent interest in peak geodetic dynamic signals (Crowell, 2021; Fang et al., 2020; Melgar et al., 2015; Ruhl et al., 2019) prompted use of unfiltered timeseries to capture peak signals for magnitude scaling laws and ground motion intensity measurements (Figures 1b and 1c). These epoch-wise threshold detection methods (i.e., Crowell et al., 2009; Dittmann et al., 2022; Hohensinn & Geiger, 2018; Hodgkinson et al., 2020; Psimoulis et al., 2018) use instantaneous measurements to estimate motion onset and therefore are a step-forward for inclusion of GNSS-seismology waveforms, but have limited “real-world” testing and most importantly mitigate false alerts for operational systems by correlating detections with proximal stations within networks or windowing in time from seismic triggers. These processes reduce the usefulness of these measurements for rapid, stand-alone decision criteria. The evolution of these detection methods has been vital for the vanguard of GNSS-based seismology, but fall short for real-time operational hazard systems to ingest the full temporal and

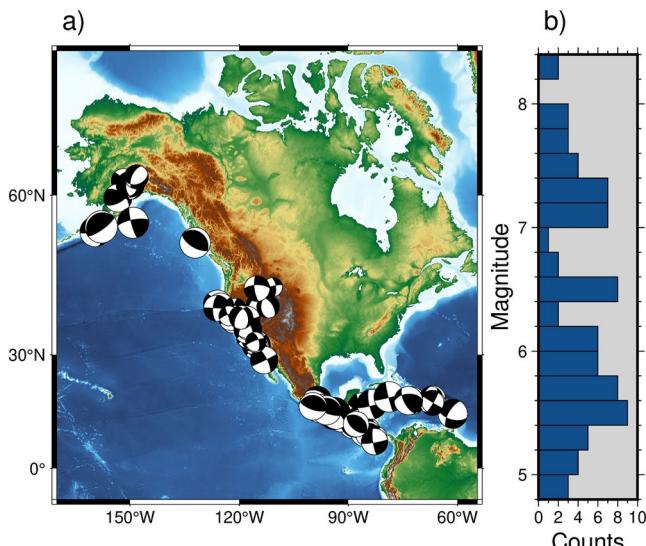


Figure 2. Map of seismic focal mechanisms used in this work and distribution of 77 event magnitudes. The number of stations used in each event is a function of the ground station network density and the magnitude-dependent sensitivity radii.

a substantial data set from which to train, validate and test the model. For optimal results, this data set requires broad spectrum noise and signal samples such that the model can “learn” and generalize our classification and distinguish signal from noise. We assembled a catalog of 1,701 station-event pairs from 77 events by cross referencing available 5 Hz GNSS observational data in the UNAVCO geodetic archive with Advanced National Seismic System Comprehensive Earthquake Catalog (COMCAT) of earthquakes greater than M_w 4.5. While 1 Hz data is more readily available, 1 Hz observable decimation undersamples certain event velocity spectra (Joyner, 1984). We observe this effect in reduced velocity amplitudes from 1 Hz data when compared to 5 Hz observables in several nearfield TDCP velocity timeseries, such as the M_w 6.2 2021 Petrolia event. For larger magnitude events it’s likely that sampling closer to 10 Hz is necessary to avoid aliasing (Shu et al., 2018), but we balance this design parameter with the need for sufficiently large available datasets for training. We assigned a conservative radius of detection for each event using ambient noise estimation from Dittmann et al. (2022). For each station-event pair within this spatial footprint, a time series window began 2 min prior to earthquake origin time (OT), and extends out in time as a function of radius (Figure 2). We conservatively buffered the radius and time window to mitigate limiting this result from the existing model. We also processed the available 5 Hz observables for a 30 min window in the hour prior to event times of the event catalog from 2017 to 2021. This noise catalog consisted of 1,507 unfiltered station-noise timeseries from 904 unique stations across a range of receiver types, geographic locations, antenna environments and atmospheric conditions, among other potential TDCP noise variance sources. Inclusion of this extended, real-world noise data set in training and validating will improve the model’s generalization, or performance on unseen data.

Current use of GNSS-derived seismic ground motion for operational EEW (Murray et al., 2018) use precise point positioning (PPP) derived topocentric coordinates to capture dynamic waveforms or static offsets relative to a stations a priori position. Instead, we align synchronous carrier phase epoch-wise changes, predicted satellite orbital velocity and line-of-sight geometry to accumulate coherent energy with respect to the shared receiver clock drift rate and directional velocities in a local reference frame. Variations of this geodetic processing method, known as time differenced carrier phase (TDCP) (van Graas & Soloviev, 2004) or variometric velocities, can record co-seismic velocity waveforms (Crowell, 2021; Grapenthin et al., 2018; Hohensinn & Geiger, 2018) as well as integrated over time into seismic displacement waveforms (Branzanti et al., 2013; Colosimo et al., 2011; Fratarcangeli et al., 2018). We processed these 5 hz measurements with the open-source SNIVEL package

frequency range of these valuable measurements into models with minimal stand-alone false alerting.

In this work, we evaluate whether existing GNSS hardware can: more reliably detect motion signals that are (a) constellating near the ambient temporal noise floor (b) with minimal false alerting (c) in a low-latency, stand-alone mode and (d) with no specific fault or network geometry. We trained a machine learning classifier on a supervised data set of GNSS velocity time series concurrent in space and time with known seismic source signals. We assembled, processed and labeled a data set of 1,701 earthquake-station high rate (5 Hz) time series pairs. We also include a substantial seismic event-free noise data set to improve model generalization. We optimized the classifier on these combined datasets with applied domain knowledge to feature selection and feature engineering that is able to combine time and frequency domain information. We present the superior performance of this classifier relative to existing methods within this motivational context. We offer advantages and implications of deploying this processing and trained model at scale for network wide monitoring, with particular emphasis on the improved sensitivity and integrity of stand-alone GNSS seismic event detection without external inputs.

2. Methods

2.1. Signals of Interest

We define our detection domain as a binary *event* or *no event* state classification. A critical component of developing a robust classification model is

Table 1

Distribution of Classification Sample Labels Used in Training/Testing Datasets by Component and Label

	East (n = 135,671)	North (n = 135,671)	Up (n = 135,671)
Non-event	94%	94%	99%
Event	5%	5%	1%
Maybe ^a	1%	1%	<1%

Note. For more information regarding the distribution of peak values, see Figure 5c.

^aMaybe's excluded from training/testing.

(Crowell, 2021) using broadcast ephemeris and narrow lane phase combinations. We chose TDCP over PPP because it is more sensitive to motion (Dittmann et al., 2022; Fang et al., 2020), and it is “lightweight” in that it does not require sophisticated corrections and is computationally inexpensive. From a machine learning perspective, this could be considered a first step in our feature engineering, or applying domain knowledge to extracting features that are correlated with motion in observed carrier phase measurements.

2.2. Feature Engineering Pipeline

Data-driven supervised machine learning models are widely used in computer vision and natural language processing due to their superior accuracy for challenging classification, regression and clustering problems. Earth scientists have adopted many of these models for geoscience research (Kong et al., 2019). Recent catalogs of historic seismic data training sets (e.g., Stanford Earthquake Data Set (Mousavi et al., 2019), INSTANCE (Michelini et al., 2021)) have contributed to benchmarking improvements of earthquake detection, phase picking, localization, and magnitude estimation (e.g. Meier et al. (2019); Mousavi et al. (2020); Kong et al. (2019)). These extensive labeled data sets enable sophisticated data-driven classifiers and deep learning models using inertial seismic data. Several geodetic applications of machine learning algorithms have demonstrated promising results with respect to seismic processes. Crocetti et al. (2021) used a random forest classifier for antenna offset detection, including due to earthquake offsets, from low-rate, 24-hr position solutions. Habboub et al. (2020) applied a neural network to coordinate time series anomaly detection applicable to specific regional datasets well above the noise floor. Dybing et al. (2021) used neural networks for earthquake detection and Lin et al. (2021) employed deep learning used for rapid event magnitude estimation; both of these studies used extensive synthetic displacement waveforms derived from real-world fault geometries and real-world PPP noise models.

In our study, we used a random forest algorithm for our classifier (Breiman, 2001) of GNSS velocities. Random forest is an ensemble of decision trees; a single decision tree is a classifier where input features are split along thresholds to separate source, or root, data from end node classifications, or leaves. An ensemble or forest of trees each vote on the feature decision criteria to select the optimal decisions toward minimizing correlated noise. Due to the infrequent nature of larger magnitude earthquakes, the event classes are naturally imbalanced but by pre-selecting specific time series of events, we have reduced this imbalance for training (Table 1) and testing. Random forest hyperparameters were selected using a grid search over the number of decision trees used, the maximum decision splits within a tree, and imbalance classification weighting strategies.

SNIVEL TDCP processing generates 5 Hz time series of the three topocentric velocity components and the clock drift rates. From these event-station pair time series of velocities, we generated feature sets to label for our supervised classification (Figure 3). Our feature samples consisted of three directional components of 30 s windows overlapping every 10 s; within these windows we included the four maximum component norm window values, window median, window median absolute deviation and window power spectral densities from the lowest frequencies bins containing periods 1–30 s as features. These features and windowing allowed our model to incorporate signal and noise amplitude in the time domain, akin to the traditional threshold approach, as well as power spectra in the frequency domain. In our binary classification, an *event* is seismic ground motion in an individual component. Labels were assigned through visual inspection as *not event* or 0, *event* or 1, and *maybe* for windows

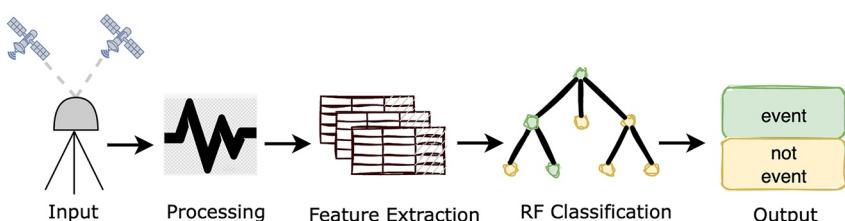


Figure 3. Schematic of our classification workflow: Inputs were 5 Hz GPS phase measurements and broadcast ephemeris, which are processed using narrow lane combinations using SNIVEL. Target labeling combined with feature extraction were used for training a supervised random forest classification model to predict event classification on testing subsets.

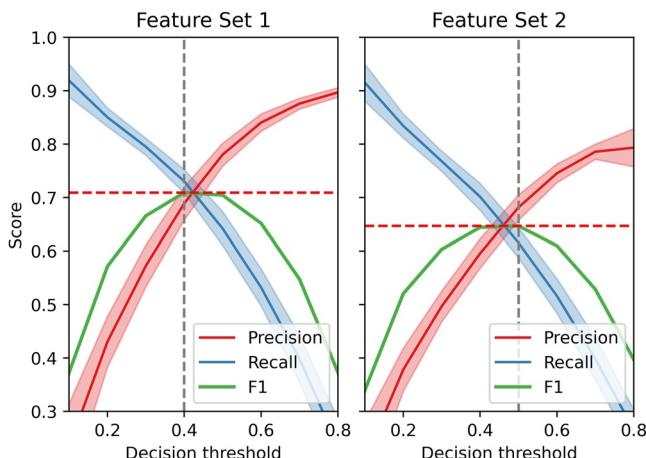


Figure 4. Mean precision, recall and F1 as a function of decision thresholds for the 10 fold nested cross validation evaluation. The shaded regions are the standard deviations across the 10 folds as a function of threshold. The dashed vertical lines are the maximum F1 decision threshold, with the dashed horizontal lines being the corresponding maximum F1 score.

events we avoided “leakage” of information from our training into our testing, including correlation of seismic waveforms from any given event observed across a network. By cross validating over 10 folds we minimized biasing our result by the relatively small testing subsets of events, and can quantify the ability of our classification model to generalize for future events. Each event was observed by a different number of stations depending on network density and sensing radius, and each station-event pair had differing number of time samples; consequently the feature vectors of training and testing were not precisely 90/10 split in samples. In each fold, we held the test set aside as “unseen,” and tuned our model using K-fold cross validation (Bishop & Nasrabadi, 2007) on the remaining training set. We implemented five inner folds in our K-fold cross validation to find the best hyperparameters. This cross validation approach allowed us to minimize overfitting the training data set and evaluate the performance of our model on unseen data as though it were running such a classifier on yet-to-occur events.

The traditional “accuracy” metric, or the ratio of the correctly classified labels relative to the total number of labels, of our classification will be less sensitive regardless of optimization choices due to the infrequent events of our imbalanced classification. Instead, we optimized on metrics that reflect accurately classifying the infrequent events. Precision, or positive predictive value, is equal to the number of true positives (TP) over the sum of TP and false positives (FP).

$$Precision = \frac{TP}{TP + FP} \quad (1)$$

recall, or sensitivity, is the number of TP over the sum of TP and false negatives (FN).

$$Recall = \frac{TP}{TP + FN} \quad (2)$$

F1 is the harmonic mean of precision and recall:

$$F_1 = 2 \times \frac{precision \times recall}{precision + recall} \quad (3)$$

Here, positive denotes motion and vice versa.

Precision and recall are approximately inversely related and each is a function of our random forest decision threshold (Figure 4). Quantifying missed detections and false alert rates is imperative for the effectiveness of any EEW system (Minson et al., 2019). We optimized hyperparameters on F1 scores, a balance of precision and

that we are not able to distinguish and excluded from testing and training. Each directional component was labeled independently. This resulted in 140,334 labels for the on-average 30 samples for 1,701 station event pairs of three component velocity time series. The event-free, noise data set included an additional 266,739 labels of three component, non-overlapping velocity time series.

We evaluated two feature extraction models. Feature set #1 was a combined array of all three directional components with a single label at each window. The horizontally concatenated components resulted in $3 \times m$ features and n samples, where m is the number of features per component ($m = 36$ in our pipeline) and n is the number of window samples. If any component was labeled “1” for event, the feature set #1 sample label was “1” for event. If a *maybe* label was present without *yes* events on the other concurrent components, the window was excluded from training/testing. Feature set two included a target vector for each component but excluded the noisier vertical signals. These vertically concatenated components resulted in m features and $2 \times n$ samples. In this extraction case any *maybe* labels were excluded from training and testing.

We employed a nested cross validation approach for unbiased testing of our data set. We initialized 10 different folds of randomly splitting the 77 events and noise catalog samples into 90% training and 10% testing. By splitting on

Table 2

10-Fold Nested Cross Validation Results Comparing Feature Set 1 Is Where All 3 Components Are Combined For Each Window, And Feature Set 2 Is Where Each Horizontal Component Is Tested Independently

		Feature set #1	Feature set #2
Precision	Mean	0.72	0.64
	stdev	0.19	0.19
Recall	Mean	0.70	0.65
	stdev	0.10	0.15
F1	Mean	0.70	0.63
	stdev	0.12	0.14

recall, but this parameter is a knob available to tune depending on societal missed detection of false alerting tolerances of a future operational system.

3. Results and Discussion

We evaluate the two optimal feature selection strategies and a range of random forest hyperparameters using a grid search. Given the F1 scores of our 10 fold nested cross validation approach (Table 2), our optimal model used feature set #1, with all available spatial components with a single target label to accumulate as much signal as possible toward our binary classification. Each train/test fold selected different optimal hyperparameter combinations for optimizing F1 scores, but the majority selected 100–200 decision trees, 100 decision splits and no class weighting with a decision threshold of 0.4 (Figure 4). This decision threshold was selected inside the cross validation of each split and applied to testing sets along with the other

hyperparameters selected. Our mean and one standard deviation nested cross validation F1 score of 0.70 ± 0.12 indicates our ability to successfully train a model using random forest. The variance in our results as a justifies our nested cross validation approach to quantify the variability in results as a function of the testing set; presumably some variability will resolve with expanded target catalogs.

3.1. Feature Importance

A benefit of random forest is that individual feature importance is readily extracted from the trained model. When evaluating feature set one, we find several aspects of the feature importances that align with our domain knowledge and therefore contribute to the explainability of our trained model. The horizontal velocity components dominate the contribution to the model (Figure 5a). GNSS ambient noise on the vertical component is much higher than that of the horizontal components and vertical seismic signal amplitudes are diminished relative to horizontal motion along horizontal strike-slip fault mechanics that are common in the spatial region of this study. These less frequent signals amidst a higher relative noise floor were harder to detect and thus contributed less to the empirical classification model. Within a horizontal component, the lower frequency spectral features had the most influence (Figure 5b). The most important frequency bins were between 6 and 15 s periods, aligned with the dominant frequencies of seismic surface waves. Our 5 Hz sampling, as compared to lower rates, boosted the detectability around the noise floor, and avoided corner frequency aliasing of certain magnitudes.

The time domain features contributed to the model, albeit much less than the lower frequency spectral content and with a more complex relationship. Figure 5c shows increasing F1 score with increasing peak velocity up until approximately a peak velocity of 25 cm/s in the east, followed by diminishing performance. We infer this to be the result of readily visibly identifiable signal events experiencing strong to very-strong shaking around 5–20 cm/s (Worden & Wald, 2020), well above the median noise floor. Infrequent, highest peak velocities (≥ 25 cm/s) might either be the result of the largest events or noise disturbances; the latter are likely degrading the performance within these peak velocity bins. Figure 5d presents a more straightforward feature relationship in the frequency features, where the greater the accumulated power in the frequency bands of greatest importance (b), the higher the performance metrics (F1, recall, precision). After an initial evaluation, we removed the highest frequency power spectral densities from our features; these are logically “noise” in our classification and not contributing. Altogether, these feature importances illustrate a key attribute of such a machine learning approach: combining features in an explainable way into an effective decision process.

3.2. Comparison With Existing Methods

A critical performance indicator is evaluating how our classification model performs over a range of test events relative to existing threshold approaches. Logic was applied to map existing continuous epoch-wise time domain threshold detection to our 30 s overlapping window target labels. For a threshold method comparison similar to the approach of Hodgkinson et al. (2020) and Dittmann et al. (2022), we estimated the noise threshold in the 2 min window prior to seismic origin time. Hodgkinson et al. (2020) characterized the stand-alone sensitivity of detection using ambient noise antecedent to an event as a Gaussian heuristic threshold. Dittmann et al. (2022) approx-

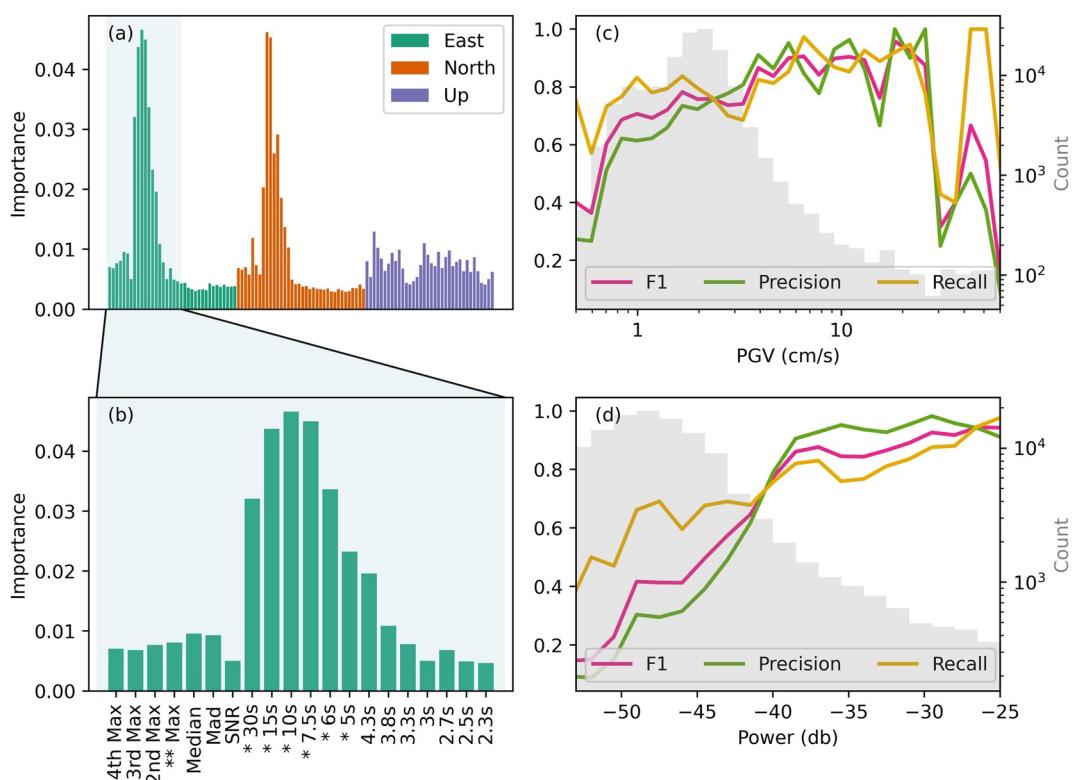


Figure 5. Panel (a) are the distribution of the feature importances across the horizontally concatenated, three spatial components feature set #1 testing. Panel (b) is a close up of the east component, with the features labeled across the x axis for closer inspection. From the left, the first six of each component are time domain features (max, min, mad) within the 30 s windows; the next 15 are the power from given frequency bins of the periodogram of the 30 s 5 Hz data, increasing in frequency from left to right, with the periods indicated. Panel (c) are precision, recall and F1 scores binned by peak velocity of each sample's east component (denoted in (b) with **). The gray shading are the counts of samples falling within the bins. Panel (d) is a similar performance measure to (c) but binned by accumulated power in the lowest frequency bins that had the highest feature importance in the model, 5–30 s period (denoted in (b) with *).

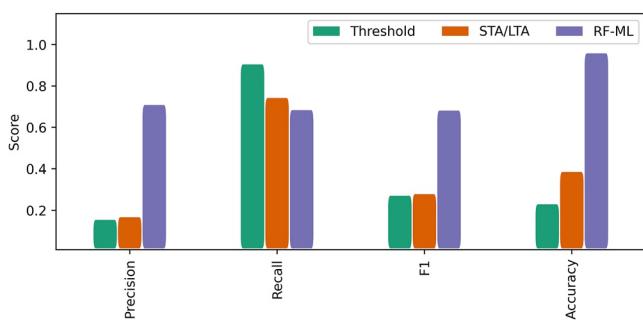


Figure 6. Performance metrics for three methods in stand-alone mode without external triggers or correlation. Threshold is the NCX2-995 approach used by Dittmann et al. (2022) that thresholds the noise based upon the 0.995 significance of a non-central chi-square distribution of the ambient noise. STA/LTA is based on Psimoulis et al. (2018) global navigation satellite system (GNSS) motion detection modified STA/LTA algorithm. RF-ML is the method presented in the work here. Optimizing on F1 in this study allows us to balance missed detections (recall) with false alerts (precision); given the amount of false alerts of the Threshold and STA/LTA, the higher recall score could be a result of regular noise triggering events.

imated the 2 min window of ground velocities as a non central chi-square (NCX2) distribution with three degrees of freedom, and then set the 0.995 confidence level value of this distribution as a noise floor approximation. Any three dimensional GV magnitude above this noise threshold after this window is considered an event, and evaluated on whether it falls within a window labeled motion or not. RT-Shake (Psimoulis et al., 2018) evolved the previous geodetic STA/LTA algorithms (Allen & Ziv, 2011; Ohta et al., 2012) by differencing instantaneous measurements from 80 epoch moving averages and then related these values to a moving window noise threshold estimate set to three times the standard error of the previous 80 epochs. This method was run on each component independently, with a single Boolean for the presence of motion on any component, and each sample window assigned a boolean based on the presence of any motion. The Dittmann et al. (2022) implementation of the threshold window in time was based upon S-Wave speeds (Crowell et al., 2013), and Psimoulis et al. (2018) modified STA/LTA correlated with surrounding stations to minimize false alerts; we did not add this logic so that we could simulate running as a stand-alone instrument.

The mean precision, F1 and accuracy from our 10 fold test of our random forest classifier outperforms the existing threshold approaches (Figure 6). In the threshold approach, recall is higher than the random forest classifier; given the large number of FP that this method triggers, we believe this value

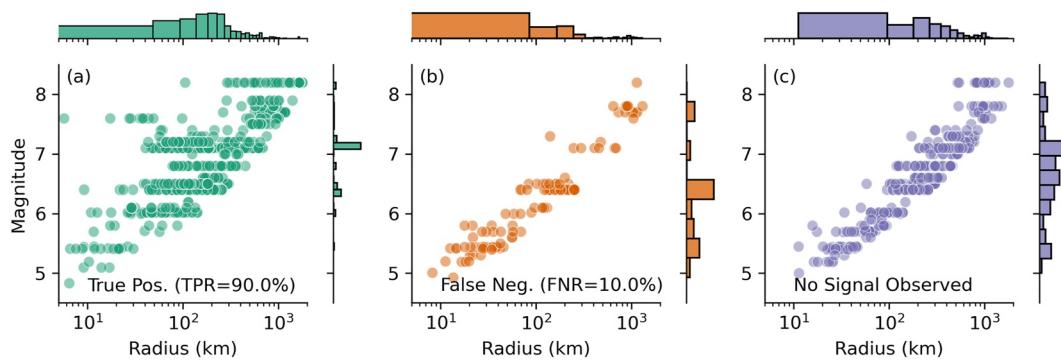


Figure 7. Performance of Random forest model developed in the work here across the entire event catalog. We reduce detection of events to a single binary for the figure. In this, each event is evaluated in a “test” split during the nested validation pipeline. This approach ensures each result depicted was evaluated as “unseen” relative to the best fit model from the training subset, and therefore representative of our model’s future performance.

is boosted by chance noise triggers occurring in windows of true motion triggering the motion boolean. This further demonstrates the value of optimizing on F1 as a balance of precision and recall to reduce biasing one decision criteria. Precision is low for both the threshold method and the STA/LTA, but for different reasons; while the precision values (Equation 1) are nearly identical, the threshold method suffers from a relatively high amount of FP, whereas the STA/LTA method low score is due to a lower amount of TP. This discrepancy is evident in the accuracy scores, where the STA/LTA outperforms the threshold approach. False positives would be decreased if using additional external information as their authors’ suggest, such as stricter time window approaches and correlating in space within networks. Such an approach would also likely improve the random forest classifiers performance but limit the utility of a stand-alone detection node. Spatio-temporal information could be incorporated into future network decision criteria.

3.3. Edge Sensitivity Detection

Detecting the largest amplitude velocity waveforms relative to ambient noise does not present a significant challenge outside of mitigating false alerting from sporadic outliers (Figure 7a), with a 98% true positive rate of events greater than M_w 6.0 and less than 100 km radius. The random forest classifier’s balance of improved false alerting relative to thresholds and improved sensitivity relative to the STA/LTA is evident for these highest seismic risks. To further investigate the random forest model performance we evaluate detecting signals closer to the noise floor. For simplicity, we bin seismic motion edge case detection into two distinct classes in what is a continuous distribution: large magnitude event seismic motion detection in the far field, and smaller magnitude events detected in the nearfield.

In the relative nearfield, much of the seismic energy passes through a station in shorter duration, varied frequency signals. Earthquake focal depth and fault slip distribution in time and space can significantly vary these waveforms as observed. Critically, the waveform signatures can appear similar to those of non geophysical processing outliers which we wish to ignore for this classification. Most existing STA/LTA methods filter these noise signals but also these valuable higher frequency dynamics. In the previous threshold methods, detection of these edge cases was a function of the ambient noise level, with low precision resulting (Figure 6) as a result of a high false positive rate. Our classifier has far less false alerts than the threshold approach in these signals, but nevertheless still presents the hardest detection domain for our classifier, evident in the missed detections of Figure 7b of events less than M_w 6.0. The left Panels of Figure 8 is an example of a smaller magnitude event (M_w 5.4) in the relative nearfield (21 km). The top four Panels ([a:d]-0) on the left of Figure 8) demonstrate that accurately detecting such an event using the threshold or modified STA/LTA approach is difficult; not only does the true signal barely exceed the noise floor, but there are numerous false alerts using both methods. The random forest classifier captures the majority of labeled motion window in addition to “ignoring” the spurious disturbance around 100 s OT that triggers all other methods evaluated 8 (e-0).

The sensitivity of GNSS to longer period surface waves are apparent at relatively great radii in the 5 hz TDCP velocity time series (Figure 7). The model detects teleseismic surface waves in unfiltered GNSS velocities from

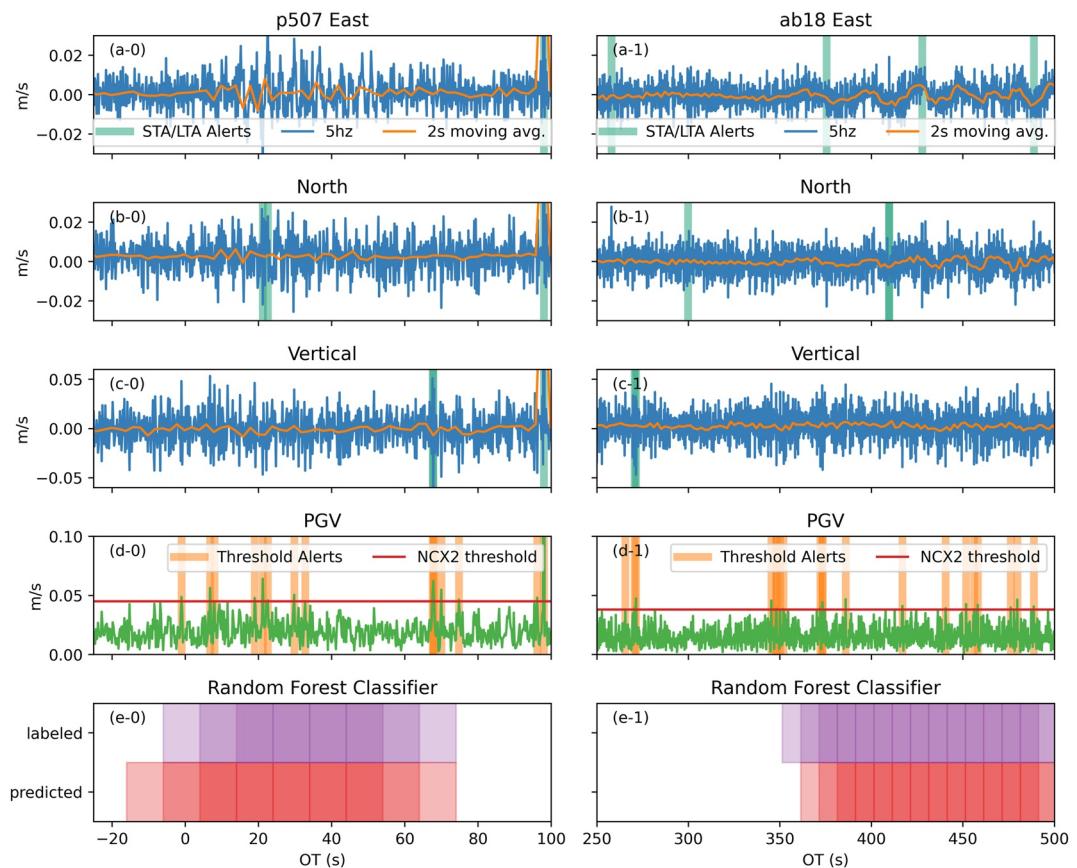


Figure 8. Velocity and detection time series two stations observing different events. The left column is from P507 observing a $M_{\text{w}} 4.1$ at 21 km epicentral radius; the right is from AB18 observing a $M_{\text{w}} 7.9$ from $\sim 1,400$ km epicentral radius. In the top velocity component Panels (a–c), we include a downsampled running mean so that the reader may readily visualize the lower frequency surface waves passing through. The teal vertical lines are alerts from the STA/LTA classifier (Psimoulis et al., 2018) on each component. Panel (d) green timeseries is the three component ground velocity; the red horizontal line is the sensitivity threshold of a 0.995 non central chi-square (ncx2) noise model (Dittmann et al., 2022), with orange vertical lines indicating a potential alert where GV greater than the threshold. (e) Panels are a comparison of the labeled feature set 1 for these event-station pairs in purple, and the results of the model prediction in red. Shading is used to distinguish overlapping windows. This event-station pair prediction is extracted from the test or unseen event collection.

a $M_{\text{w}} 8.2$ (USGS event ID: ak0219neiszm) at 1,780 km epicentral radius in real-time with no external corrections; the right-hand Panels of Figure 8 provides an example of this detection. Future analysis could investigate the range of geodetic teleseism detection with respect to larger magnitude event directivity, attenuation and observational networks. In Figure 8d-1, the amplitude of the ground velocity magnitude of these long period signals is insufficient to cross the traditional noise threshold with consistency, and there are many antecedent false alerts. The modified STA/LTA RT-Shake approach does not identify the majority of the long period waves either (Figures 8a and 8c-1), while the random forest classifier in the bottom Panel only misses the first window (Figure 8e-1).

3.4. Decision Latency

Delay in alerting is critical to EEW. While our model is trained, tested, and validated on overlapping windows every 10 s, we evaluate running the model at once per second, the current US EEW (Murray et al., 2018) geodetic input rate (Figure 9). On testing data not used in model training, we find a delay relative to the estimated P -wave, $\sim 3\text{--}5$ s under 15 km exists in the current approach. Coarse P - and S -wave arrivals are estimated using the iasp91 model (Kennett & Engdahl, 1991); future work more accurately quantifying these phase arrivals such as the approach of Goldberg et al. (2018) would not only more accurately represent timing performance but also useful for training more sophisticated ground motion models. GNSS velocities using this current approach cannot

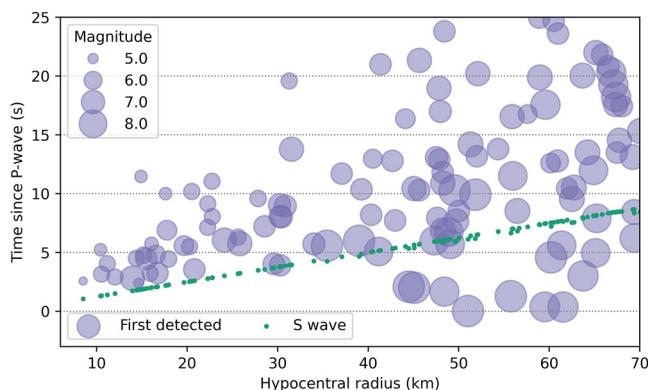


Figure 9. Time of first detection of all individual event-station pairs within 70 km radius relative to estimated P -wave arrival time (iasp91 travel time model) as a function of radius. Green dots are the estimated S -wave arrivals at the event-station pairs used in this study shown for reference. Purple circles are centered on the time of first detection after the origin time (OT), where the diameter is scaled to the event magnitude. These results are from the classifier run at 1 Hz on unseen testing sets to simulate a real-time operational mode.

cross-validation pipeline (Section 2.1) to select hyperparameters for training a complete model on all available labeled data for future “unseen” events. Such “unseen” events include this set-aside noise testing set. We confirmed there were no concurrent events greater than M_w 4.0 in the USGS COMCAT catalog within the relevant spatial footprint and all other sources of noise or disturbances (signal multipath, oscillators, atmospheric anomalies, etc) remained in the test set. We assigned labels of *non event* to all target vectors associated with feature extraction. This allowed us to quantify ambient noise performance, or false alarm rate (Figure 10a) using the detection methods previously described in Section 3.2 from 860 unique stations from Alaska to the Caribbean across a range of potential TDCP noise or disturbance sources.

The random forest classifier was less susceptible to false alerts over the window tested than the threshold and STA/LTA approaches. The two threshold models have the highest rates to false alerting, an anticipated result based upon the precision metric reported in Figure 6. Station variations present in the random forest approach (Figure 10b) suggest the current random forest model has some station or time noise dependence not correlated with the variations of other detection methods. Future inclusion of more extensive noise training datasets into our detection classifier and possibly data augmentation techniques would likely be beneficial toward training on the widest variety of noise scenarios and optimizing feature engineering for these complex noise environments.

4. Conclusion

We applied an existing machine learning algorithm and sample splitting pipeline techniques to training, validating and testing a seismic motion detection classifier from 5 Hz TDCP GNSS velocities. We leveraged nearly 20 years of 5 Hz GNSS data archives for training a classification model that outperforms existing threshold approaches for detecting motion in stand-alone mode. The classifier combines time domain and frequency domain features to match the sensitivity of the threshold method without the false alerts, and matches the minimal false alerting of the STA/LTA with improved sensitivity. Given the agreement that GNSS velocities have with existing ground motion models (Crowell et al., 2022) and the increased confidence in separating signal from noise demonstrated here, these GNSS velocities can operationally contribute to ground motion measurements. The alert latency of this current model does not match the sensitivity of existing inertial infrastructure. A complementary approach using the information available at the time, including lowest latency p -wave characterization from inertial sensors and unsaturated velocity estimation from GNSS provides an optimal solution for existing dense multi-sensor networks. For less dense networks of either sensor type, it is more critical to establish a decision criteria for balancing timing, noise and accuracy of these independent observation systems. Further investigation of integrating the processing and classifying approach of this manuscript with the sensitivity of co-located MEMS sensors (Goldberg & Bock, 2017) would advantageously overlap seismic and geodetic traditional boundaries.

reliably be used for earliest phase picking, but can rapidly contribute to ground motion models or peak motion scaling laws (Fang et al., 2020). Given the feature importances of the classifier (Figure 5a), we interpret delays to be the result of the classifier trained on the relatively longer period signals visible within the noise. Depending on source magnitude and travel path, these could be P -, S -, surface-waves or some convolution of the energy traveling through the GNSS location. Variance in delays in the near field are likely due to inherent limitations of modeling rupture as a point source at proximal locations (Goldberg et al., 2021) and possibly related to errors of the iasp91 travel time model. Future work will address the possible limitations or delays introduced by our visual classification labeling. It is worth repeating that this assessment uses no external input or seismic triggering.

3.5. Ambient Noise Data Set

In addition to evaluating the performance within the bespoke event and noise data sets, we also evaluated the performance of the method during periods of quiescence to further quantify relative false alert rates. Our unseen testing set consisted of 1,321 30-min velocity timeseries from 2019 to 2021, not included with the original nested cross-validation data. We ran five-fold cross validation on the entire event and noise labeled data set from the nested

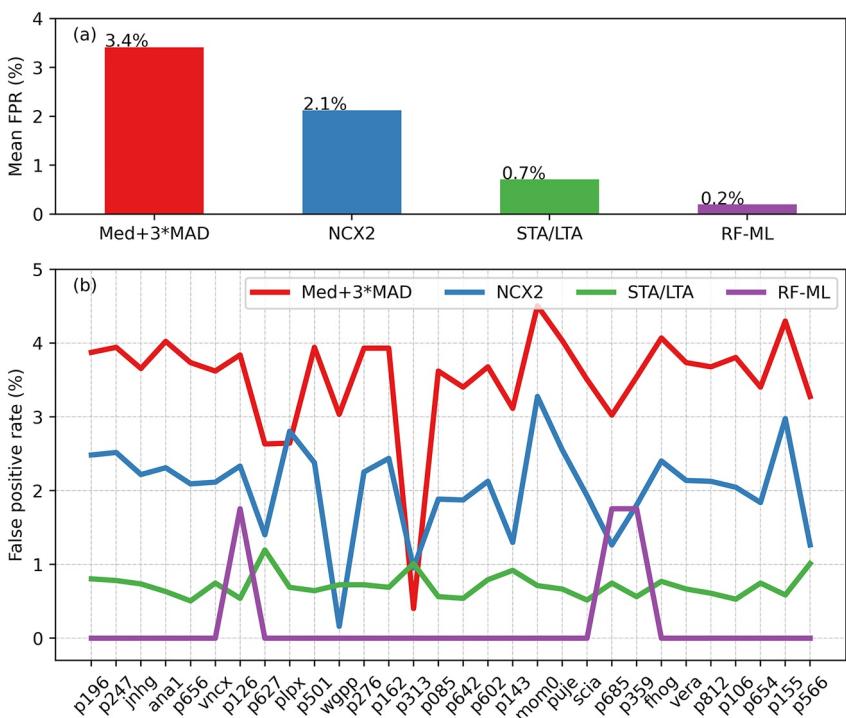


Figure 10. Panel (a) is mean false positive rates (FPR) from 1,321 spatially distributed, 30 min duration of TDCP 5 Hz velocities from windows prior to events in 2019–2021. Methods include: median plus three times the median absolute deviation threshold of Hodgkinson et al. (2020), non-central chi-square of Dittmann et al. (2022) NCX2 using alpha value of 0.995, the modified STA/LTA implemented by Psimoulis et al. (2018) and the random forest machine learning classifier developed in this work (RF-ML). Panel (b) is a distribution of each method of a randomly chosen subset of stations to illustrate some of the station variability to the reader.

Current 5 Hz GNSS observation data streams are too verbose for many bandwidth limited remote hardware; this presents an exciting opportunity for edge processing at potentially much higher rates (Shu et al., 2018), or experimental lean 5 Hz carrier phase data streams. Our method presented here does not use a sophisticated machine learning model, yet has improved detection relative to existing approaches; much improvement remains, especially with expanded datasets across global geodetic networks and/or synthetics or data augmentation for training, validation and testing of neural networks and deep learning models.

With an expanding availability and access to real-time GNSS streaming networks, the seismological community stands to benefit from this signal of opportunity for rapid ground motion detection for earthquake and tsunami source characterization. Furthermore, the vast industry of GNSS position, navigation and timing users catalyzing the expansion of these GNSS real-time networks will benefit from improved automated alerting of reference station motion onset. Future work will include integrating this classifier amongst existing and future automated GNSS carrier phase disturbance characterization methods, including space weather disturbances (Jiao et al., 2017), oscillator anomalies (Liu & Morton, 2022), radio frequency interference and signal multipath.

Data Availability Statement

The 5 Hz global navigation satellite system (GNSS) data used for time differenced carrier phase (TDCP) processing in the study are available from the Geodetic Facility for the Advancement of Geoscience (GAGE) (GNSSs) archives as maintained by UNAVCO, Inc. The data are available in RINEX (v.2.11) format at <https://data.unavco.org/archive/gnss/highrate/5-Hz/rinex/>. Earthquake depths, locations, and magnitudes came from the Advanced National Seismic System (ANSS) Comprehensive Catalog of Earthquake Events and Products (<https://earthquake.usgs.gov/data/comcat/>). Arrival times are calculated using the iasp91 velocity model as implemented by Incorporated Research Institutions for Seismology (IRIS) Web Services (<http://service.iris.edu/irisws/traveltime/>). SNIVEL code used for TDCP velocity processing is developed openly at <https://github.com/crowellbw/SNIVEL>.

(Accessed December 2021) (Crowell, 2021). SNIVEL 5 Hz velocity timeseries used in this study are preserved at <https://doi.org/10.5281/zenodo.6588601>. Version 1.0.1 of the scikit-learn software used for random forest classification is preserved at <https://doi.org/10.5281/zenodo.5596244> and developed openly at <https://github.com/scikit-learn/scikit-learn> (Pedregosa et al., 2011). Version v0.5.0 of PyGMT used for generating the map is preserved at <https://doi.org/10.5281/zenodo.5607255> and developed openly at <https://github.com/GenericMappingTools/pygmt> (Wessel et al., 2019). Figures were made with Matplotlib version 3.5.1 (Caswell et al., 2021), available under the Matplotlib license at <https://matplotlib.org/>.

Acknowledgments

We thank Brendan Crowell, an anonymous reviewer and editors Paul Tregoning and Yosuke Aoki for their thoughtful feedback. This material is based on services provided by the GAGE Facility, operated by UNAVCO, Inc., with support from the National Science Foundation, the National Aeronautics and Space Administration, and the U.S. Geological Survey under NSF Cooperative Agreement EAR-1724794. High-rate processing and machine learning for geoscience and hazards research is supported by NSF OAC-1835791.

References

Allen, R. M., & Ziv, A. (2011). Application of real-time GPS to earthquake early warning. *Geophysical Research Letters*, 38(16). <https://doi.org/10.1029/2011gl047947>

Bishop, C. M., & Nasrabadi, N. M. (2007). Pattern recognition and machine learning. *Journal of Electronic Imaging*, 16, 049901.

Bock, Y., Prawirodirdjo, L., & Melbourne, T. I. (2004). Detection of arbitrarily large dynamic ground motions with a dense high-rate GPS network. *Geophysical Research Letters*, 31(6). <https://doi.org/10.1029/2003gl019150>

Branzanti, M., Colosimo, G., Crespi, M., & Mazzoni, A. (2013). Gps near-real-time coseismic displacements for the great Tohoku-Oki earthquake. *IEEE Geoscience and Remote Sensing Letters*, 10(2), 372–376. <https://doi.org/10.1109/LGRS.2012.2207704>

Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. <https://doi.org/10.1023/a:1010933404324>

Caswell, T. A., Droettboom, M., Lee, A., de Andrade, E. S., Hoffmann, T., Hunter, J., et al. (2021). matplotlib/matplotlib: Rel: V3.5.1. *Zenodo*. <https://doi.org/10.5281/zenodo.5773480>

Colombelli, S., Allen, R. M., & Zollo, A. (2013). Application of real-time GPS to earthquake early warning in subduction and strike-slip environments. *Journal of Geophysical Research: Solid Earth*, 118(7), 3448–3461. <https://doi.org/10.1002/jgrb.50242>

Colosimo, G., Crespi, M., & Mazzoni, A. (2011). Real-time GPS seismology with a stand-alone receiver: A preliminary feasibility demonstration. *Journal of Geophysical Research*, 116(B11). <https://doi.org/10.1029/2010jb007941>

Crocetti, L., Schartner, M., & Soja, B. (2021). Discontinuity detection in GNSS station coordinate time series using machine learning. *Remote Sensing*, 13(19), 3906. <https://doi.org/10.3390/rs13193906>

Crowell, B., DeGrande, J., Dittmann, T., & Ghent, J. (2022). Validation of peak ground velocities recorded on very-highrate gnss against nga-west2 ground motion models. *Seismological Society of America Annual Meeting* 2022.

Crowell, B. W. (2021). Near-field strong ground motions from GPS-derived velocities for 2020 intermountain Western United States earthquakes. *Seismological Research Letters*, 92(2A), 840–848. <https://doi.org/10.1785/0220200325>

Crowell, B. W., Bock, Y., & Squibb, M. B. (2009). Demonstration of earthquake early warning using total displacement waveforms from real-time GPS networks. *Seismological Research Letters*, 80(5), 772–782. <https://doi.org/10.1785/gssrl.80.5.772>

Crowell, B. W., Melgar, D., Bock, Y., Haase, J. S., & Geng, J. (2013). Earthquake magnitude scaling using seismogeodetic data. *Geophysical Research Letters*, 40(23), 6089–6094. <https://doi.org/10.1002/2013gl058391>

Dittmann, T., Hodgkinson, K., Morton, J., Mencin, D., & Mattioli, G. S. (2022). Comparing sensitivities of geodetic processing methods for rapid earthquake magnitude estimation. *Seismological Research Letters*, 93(3), 1497–1509. <https://doi.org/10.1785/0220210265>

Dybing, S., Melgar, D., Thomas, A., Hodgkinson, K., & Mencin, D. (2021). Detecting earthquakes in noisy real-time gnss data with machine learning. *American Geophysical Union Fall Meeting* 2021. Retrieved from <https://agu.confex.com/agu/fm21/meetingapp.cgi/Paper/820583>

Fang, R., Zheng, J., Geng, J., Shu, Y., Shi, C., & Liu, J. (2020). Earthquake magnitude scaling using peak ground velocity derived from high-rate GNSS observations. *Seismological Research Letters*, 92(1), 227–237. <https://doi.org/10.1785/0220190347>

Fratarcangeli, F., Savastano, G., D'Achille, M. C., Mazzoni, A., Crespi, M., Riguzzi, F., et al. (2018). Vadose reliability and accuracy of real-time displacement estimation: Application to the central Italy 2016 earthquakes. *Remote Sensing*, 10(8), 1201. <https://doi.org/10.3390/rs10081201>

Goldberg, D. E., & Bock, Y. (2017). Self-contained local broadband seismogeodetic early warning system: Detection and location. *Journal of Geophysical Research: Solid Earth*, 122(4), 3197–3220. <https://doi.org/10.1002/2016jb013766>

Goldberg, D. E., Melgar, D., Bock, Y., & Allen, R. M. (2018). Geodetic observations of weak determinism in rupture evolution of large earthquakes. *Journal of Geophysical Research: Solid Earth*, 123(11), 9950–9962. <https://doi.org/10.1029/2018JB015962>

Goldberg, D. E., Melgar, D., Hayes, G. P., Crowell, B. W., & Sahakian, V. J. (2021). A ground-motion model for GNSS peak ground displacement. *Bulletin of the Seismological Society of America*, 111(5), 2393–2407. <https://doi.org/10.1785/0120210042>

Grapenthin, R., West, M., & Freymueller, J. (2017). The utility of GNSS for earthquake early warning in regions with sparse seismic Networks the utility of GNSS for earthquake early warning in regions with sparse seismic networks. *Bulletin of the Seismological Society of America*, 107(4), 1883–1890. <https://doi.org/10.1785/0120160317>

Grapenthin, R., West, M., Tape, C., Gardine, M., & Freymueller, J. (2018). Single-frequency instantaneous GNSS velocities resolve dynamic ground motion of the 2016 Mw 7.1 Iniskin, Alaska, earthquake. *Seismological Research Letters*, 89(3), 1040–1048. <https://doi.org/10.1785/0220170235>

Habboub, M., Psimoulis, P. A., Bingley, R., & Rothacher, M. (2020). A multiple algorithm approach to the analysis of GNSS coordinate time series for detecting geohazards and anomalies. *Journal of Geophysical Research: Solid Earth*, 125(2). <https://doi.org/10.1029/2019jb018104>

Hodgkinson, K. M., Mencin, D. J., Feaux, K., Sievers, C., & Mattioli, G. S. (2020). Evaluation of earthquake magnitude estimation and event detection thresholds for real-time GNSS networks: Examples from recent events captured by the network of the Americas. *Seismological Research Letters*, 91(3), 1628–1645. <https://doi.org/10.1785/0220190269>

Hohensinn, R., & Geiger, A. (2018). Stand-alone GNSS sensors as velocity seismometers: Real-time monitoring and earthquake detection. *Sensors*, 18(11), 3712. <https://doi.org/10.3390/s18113712>

Jiao, Y., Hall, J. J., & Morton, Y. T. (2017). Performance evaluation of an automatic GPS ionospheric phase scintillation detector using a machine-learning algorithm. *Navigation*, 64(3), 391–402. <https://doi.org/10.1002/navi.188>

Joyner, W. (1984). A scaling law for the spectra of large earthquakes. *Bulletin of the Seismological Society of America*, 74(4), 1167–1188.

Kawamoto, S., Hiyama, Y., Ohta, Y., & Nishimura, T. (2016). First result from the GEONET real-time analysis system (REGARD): The case of the 2016 Kumamoto earthquakes. *Earth Planets and Space*, 68(1), 190. <https://doi.org/10.1186/s40623-016-0564-4>

Kennett, B. L. N., & Engdahl, E. R. (1991). Traveltimes for global earthquake location and phase identification. *Geophysical Journal International*, 105(2), 429–465. <https://doi.org/10.1111/j.1365-246X.1991.tb06724.x>

Kong, Q., Trugman, D. T., Ross, Z. E., Bianco, M. J., Meade, B. J., & Gerstoft, P. (2019). Machine learning in seismology: Turning data into insights. *Seismological Research Letters*, 90(1), 3–14. <https://doi.org/10.1785/0220180259>

Larson, K. M. (2009). GPS seismology. *Journal of Geodesy*, 83(3–4), 227–233. <https://doi.org/10.1007/s00190-008-0233-x>

Larson, K. M., Bodin, P., & Gomberg, J. (2003). Using 1-Hz GPS data to measure deformations caused by the denali fault earthquake. *Science*, 300(5624), 1421–1424. <https://doi.org/10.1126/science.1084531>

Lin, J., Melgar, D., Thomas, A. M., & Searcy, J. (2021). Early warning for great earthquakes from characterization of crustal deformation patterns with deep learning. *Journal of Geophysical Research: Solid Earth*, 126(10). <https://doi.org/10.1029/2021jb022703>

Liu, Y., & Morton, Y. J. (2022). Improved automatic detection of GPS satellite oscillator anomaly using a machine learning algorithm. *Navigation: Journal of the Institute of Navigation*, 69(1), 500. <https://doi.org/10.33012/navi.500>

Meier, M., Ross, Z. E., Ramachandran, A., Balakrishna, A., Nair, S., Kundzicz, P., et al. (2019). Reliable real-time seismic signal/noise discrimination with machine learning. *Journal of Geophysical Research: Solid Earth*, 124(1), 788–800. <https://doi.org/10.1029/2018jb016661>

Melgar, D., Crowell, B. W., Geng, J., Allen, R. M., Bock, Y., Riquelme, S., et al. (2015). Earthquake magnitude calculation without saturation from the scaling of peak ground displacement. *Geophysical Research Letters*, 42(13), 5197–5205. <https://doi.org/10.1002/2015gl064278>

Melgar, D., Crowell, B. W., Melbourne, T. I., Szeliga, W., Santillan, M., & Scrivner, C. (2020). Noise characteristics of operational real-time high-rate GNSS positions in a large aperture network. *Journal of Geophysical Research: Solid Earth*, 125(7). <https://doi.org/10.1029/2019jb019197>

Michelini, A., Cianetti, S., Gaviano, S., Giunchi, C., Jozinović, D., & Lauciani, V. (2021). Instance – The Italian seismic dataset for machine learning. *Earth System Science Data*, 13(12), 5509–5544. <https://doi.org/10.5194/essd-13-5509-2021>

Minson, S. E., Baltay, A. S., Cochran, E. S., Hanks, T. C., Page, M. T., McBride, S. K., et al. (2019). The limits of earthquake early warning accuracy and best alerting strategy. *Scientific Reports*, 9(1), 2478. <https://doi.org/10.1038/s41598-019-39384-y>

Minson, S. E., Murray, J. R., Langbein, J. O., & Gomberg, J. S. (2014). Real-time inversions for finite fault slip models and rupture geometry based on high-rate GPS data. *Journal of Geophysical Research: Solid Earth*, 119(4), 3201–3231. <https://doi.org/10.1002/2013jb010622>

Mousavi, S. M., Ellsworth, W. L., Zhu, W., Chuang, L. Y., & Beroza, G. C. (2020). Earthquake transformer—An attentive deep-learning model for simultaneous earthquake detection and phase picking. *Nature Communications*, 11(1), 3952. <https://doi.org/10.1038/s41467-020-17591-w>

Mousavi, S. M., Sheng, Y., Zhu, W., & Beroza, G. C. (2019). STanford EArthquake dataset (STEAD): A global data set of seismic signals for ai. *IEEE Access*, 7, 179464–179476. <https://doi.org/10.1109/access.2019.2947848>

Murray, J. R., Crowell, B. W., Grapenthin, R., Hodgkinson, K., Langbein, J. O., Melbourne, T., et al. (2018). Development of a geodetic component for the U.S. West coast earthquake early warning system. *Seismological Research Letters*, 89(6), 2322–2336. <https://doi.org/10.1785/0220180162>

Nikolaidis, R. M., Bock, Y., Jonge, P. J. d., Shearer, P., Agnew, D. C., & Domselaar, M. V. (2001). Seismic wave observations with the global positioning system. *Journal of Geophysical Research*, 106(B10), 21897–21916. <https://doi.org/10.1029/2001jb000329>

Ohta, Y., Kobayashi, T., Tsushima, H., Miura, S., Hino, R., Takasu, T., et al. (2012). Quasi real-time fault model estimation for near-field tsunami forecasting based on RTK-GPS analysis: Application to the 2011 Tohoku-Oki earthquake (Mw 9.0). *Journal of Geophysical Research*, 117(B2). <https://doi.org/10.1029/2011jb008750>

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12, 2825–2830.

Psimoulis, P. A., Houlié, N., Habboub, M., Michel, C., & Rothacher, M. (2018). Detection of ground motions using high-rate GPS time-series. *Geophysical Journal International*, 214(2), 1237–1251. <https://doi.org/10.1093/gji/ggy198>

Ruhl, C. J., Melgar, D., Geng, J., Goldberg, D. E., Crowell, B. W., Allen, R. M., et al. (2019). A global database of strong-motion displacement GNSS recordings and an example application to PGD scaling. *Seismological Research Letters*, 90(1), 271–279. <https://doi.org/10.1785/0220180177>

Shu, Y., Fang, R., Li, M., Shi, C., Li, M., & Liu, J. (2018). Very high-rate GPS for measuring dynamic seismic displacements without aliasing: Performance evaluation of the variometric approach. *GPS Solutions*, 22(4), 121. <https://doi.org/10.1007/s10291-018-0785-z>

van Graas, F., & Soloviev, A. (2004). Precise velocity estimation using a stand-alone GPS receiver. *Navigation*, 51(4), 283–292. <https://doi.org/10.1002/j.2161-4296.2004.tb00359.x>

Wessel, P., Luis, J. F., Uieda, L., Scharroo, R., Wobbe, F., Smith, W. H. F., & Tian, D. (2019). The generic mapping tools version 6. *Geochemistry, Geophysics, Geosystems*, 20(11), 5556–5564. <https://doi.org/10.1029/2019GC008515>

Worden, C. B., & Wald, D. J. (2020). Shakemap manual online: Technical manual, user's guide, and software guide. *U. S. geological survey*, 1–156. Retrieved from <http://usgs.github.io/shakemap/>